From UPS to Silicon

an end-to-end evaluation of Data Center Efficiency

Michael K Patterson, PhD, PE
Annabelle Pratt, PhD

Digital Enterprise Group Architecture and Planning

Legal Disclaimer

Notice: This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information. Contact your local Intel sales office or your distributor to obtain the latest specification before placing your product order.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications, product descriptions, and plans at any time, without notice.

All products, dates, and figures are preliminary for planning purposes and are subject to change without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

64-bit Intel® XeonTM processors with Intel® EM64T requires a computer system with a processor, chipset, BIOS, OS, device drivers and applications enabled for Intel EM64T. Processor will not operate (including 32-bit operation) without an Intel EM64T-enabled BIOS. Performance will vary depending on your hardware and software configurations. Intel EM64T-enabled OS, BIOS, device drivers and applications may not be available. Check with your vendor for more information.

Irwindale and Woodcrest discussed herein may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's website at http://www.intel.com.

Introduction

Overview

•Problem Statement

•Metrics

•Data Center review and evaluation

•Power Pareto

•Opportunities

•Silicon

•Power Delivery

•Cooling

•Summary

The needs of High Density Data Centers challenge the limits

• Legacy Data Centers can not keep pace with the power required to provide the needed compute density of even today's modest users

• New Data Centers densities are pushing the limits of air-cooling and require very precise designs and operational strategies

• Reducing power and increasing efficiencies will allow greater computing performance at the same or lower TCO

Metrics

Watts/sq ft…square feet of what?

•Campus? Raised floor? Equipment footprint?

•Work cell is the smallest repeatable unit, includes cold aisle, rack, and hot aisle.

•*Advantages:* best measure of a data center's capability to handle a specific load or achieve a specific density, while removing room to room variability in non-compute raised floor area.

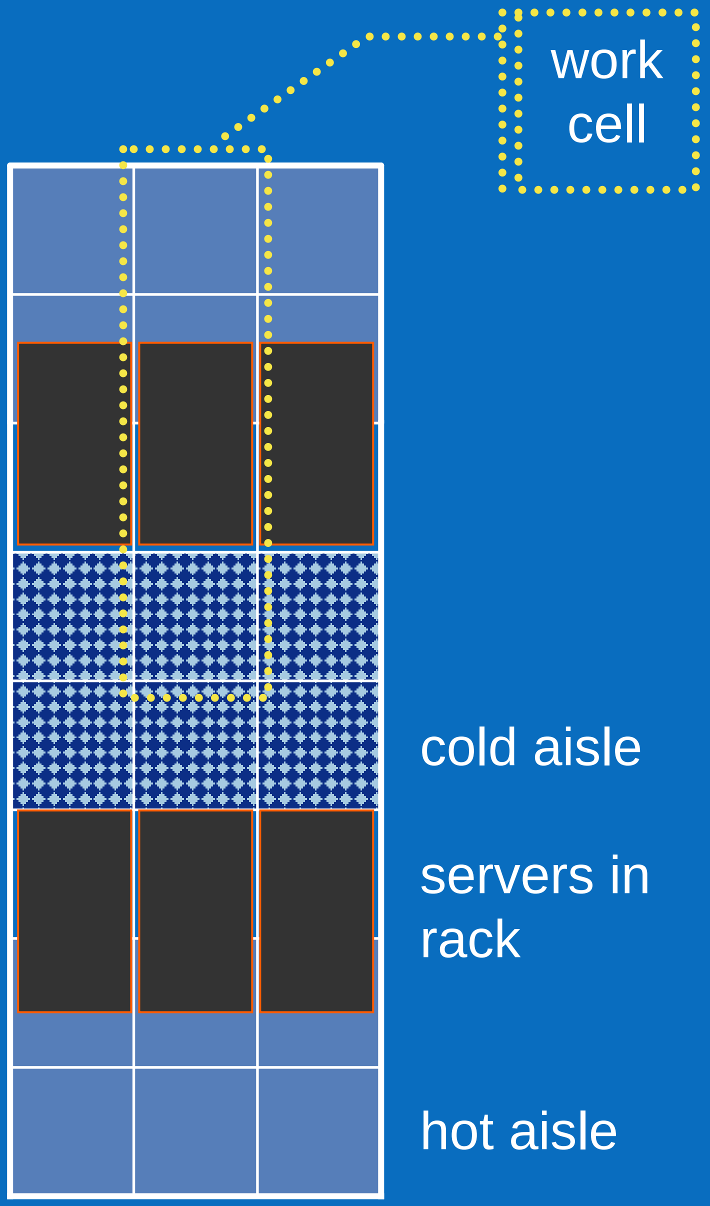

Rack density capability

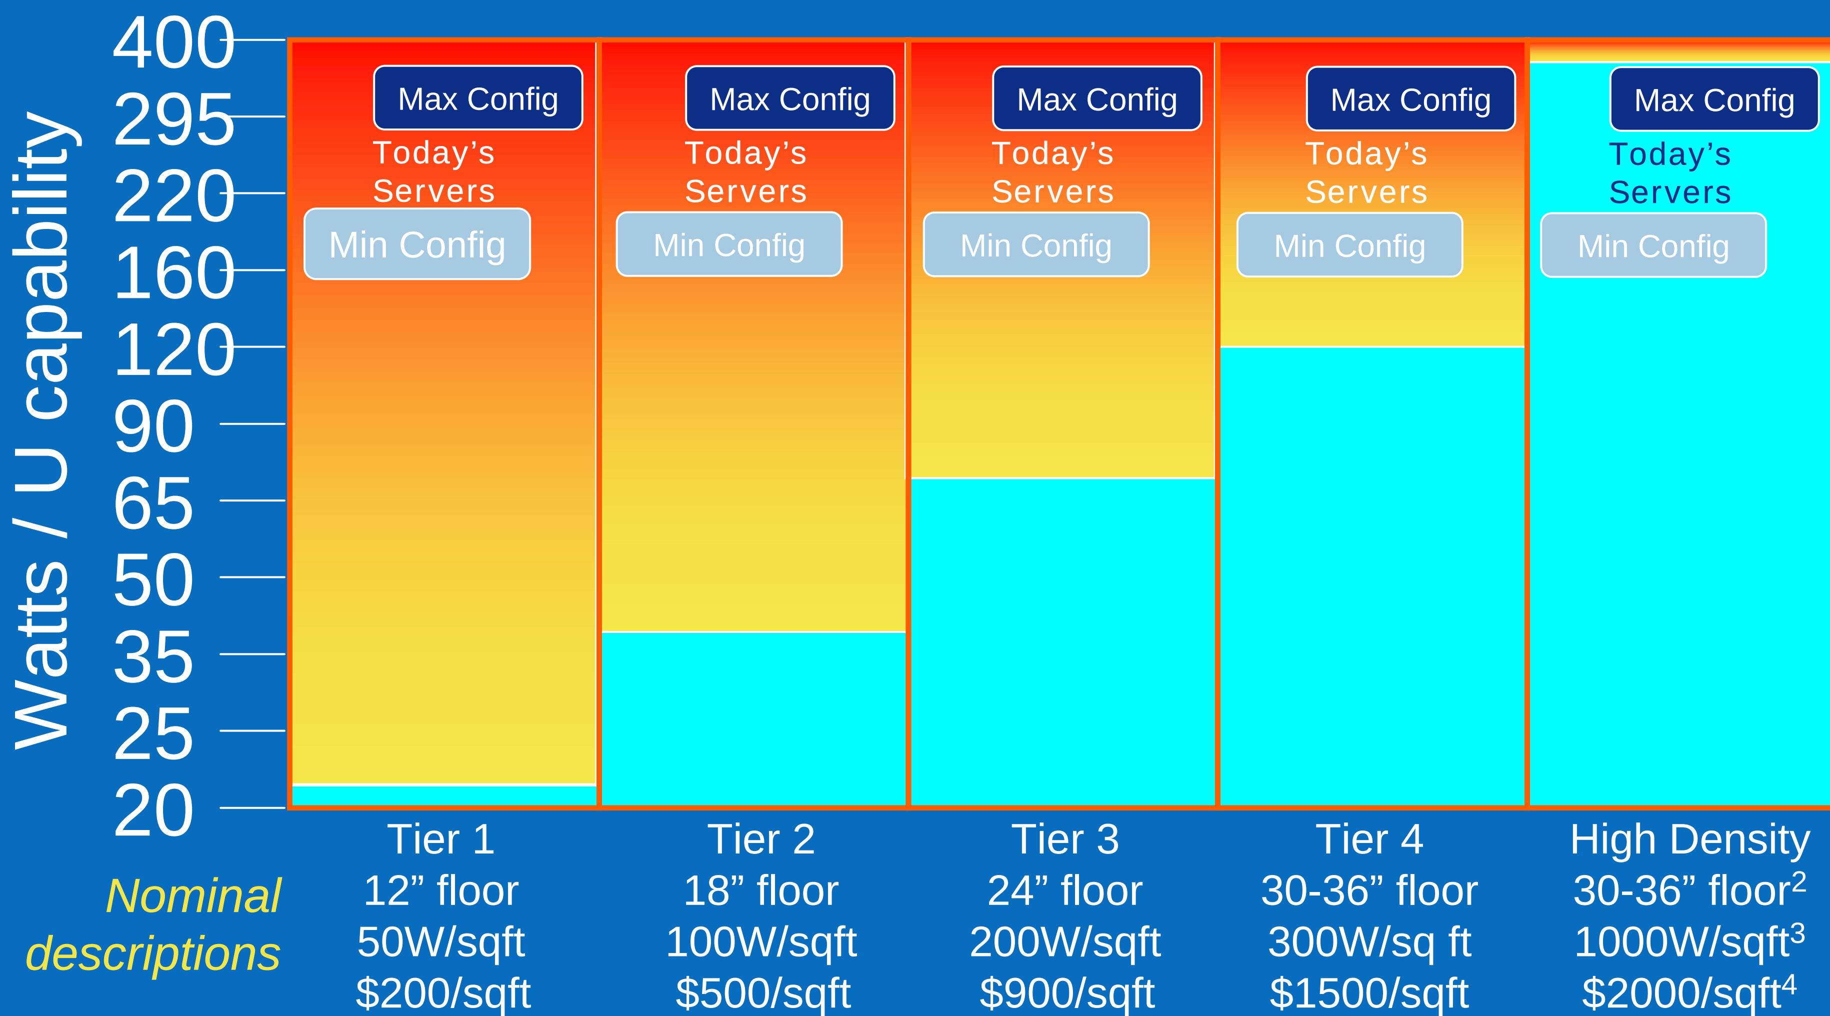

Notes: 1: per Uptime Institute defn, 2: raised floor height, 3: sq ft of work cell, 4: capital cost ROM

Subject Data Center Overview

more than 150 14 kW racks

40 1U dual processor servers per rack

UPS and redundant power feeds

875 watts/square foot of work cell

320,000 CFM of airflow

30” raised floor

power and interconnect overhead

open plenum underneath

DC Overview

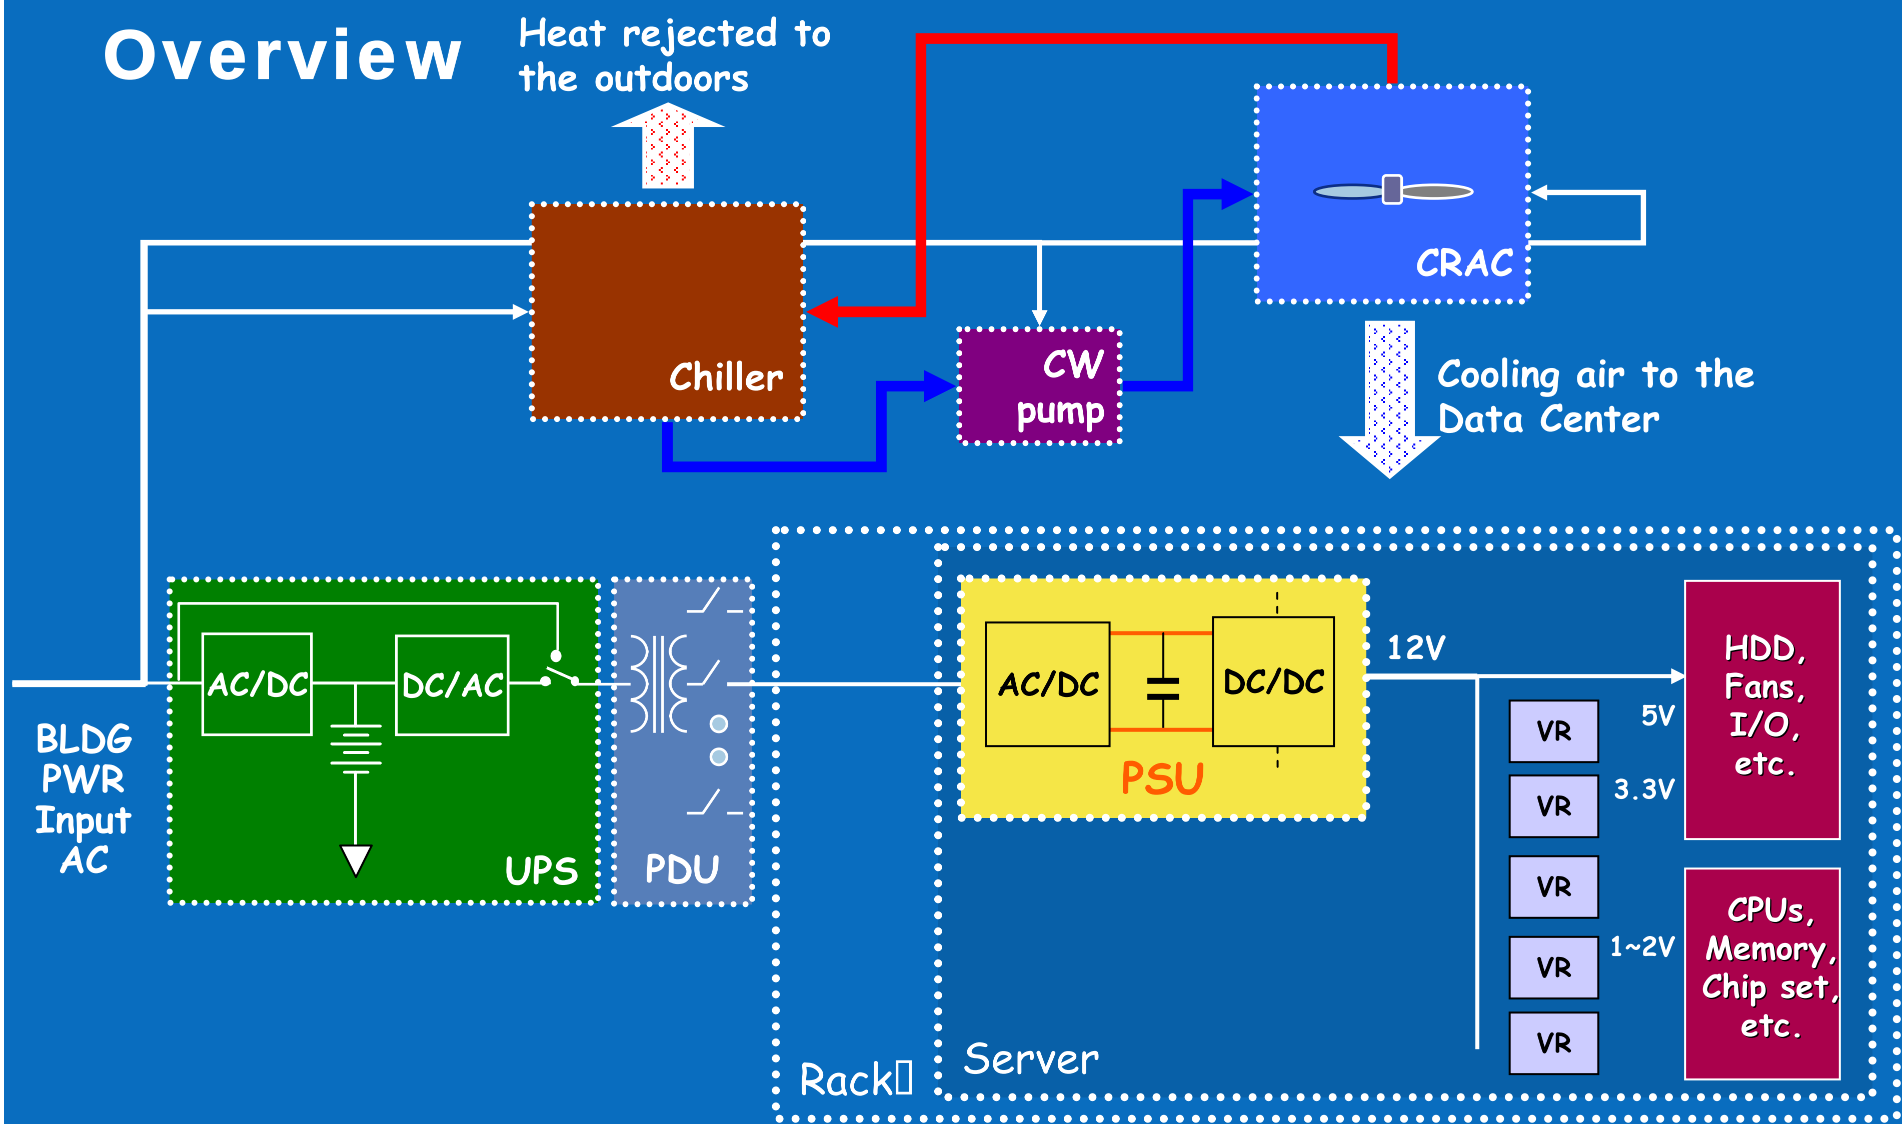
Heat rejected to the outdoors
Chiller
CW pump
CRAC
Cooling air to the Data Center
BLDG PWR Input AC
AC/DC
DC/AC
UPS
PDU
AC/DC
DC/DC
PSU
12V
VR
VR
VR
VR
VR
5V
3.3V
1~2V
HDD, Fans, I/O, etc.
CPUs, Memory, Chip set, etc.
Rack
Server

Server / Data Center power use – today

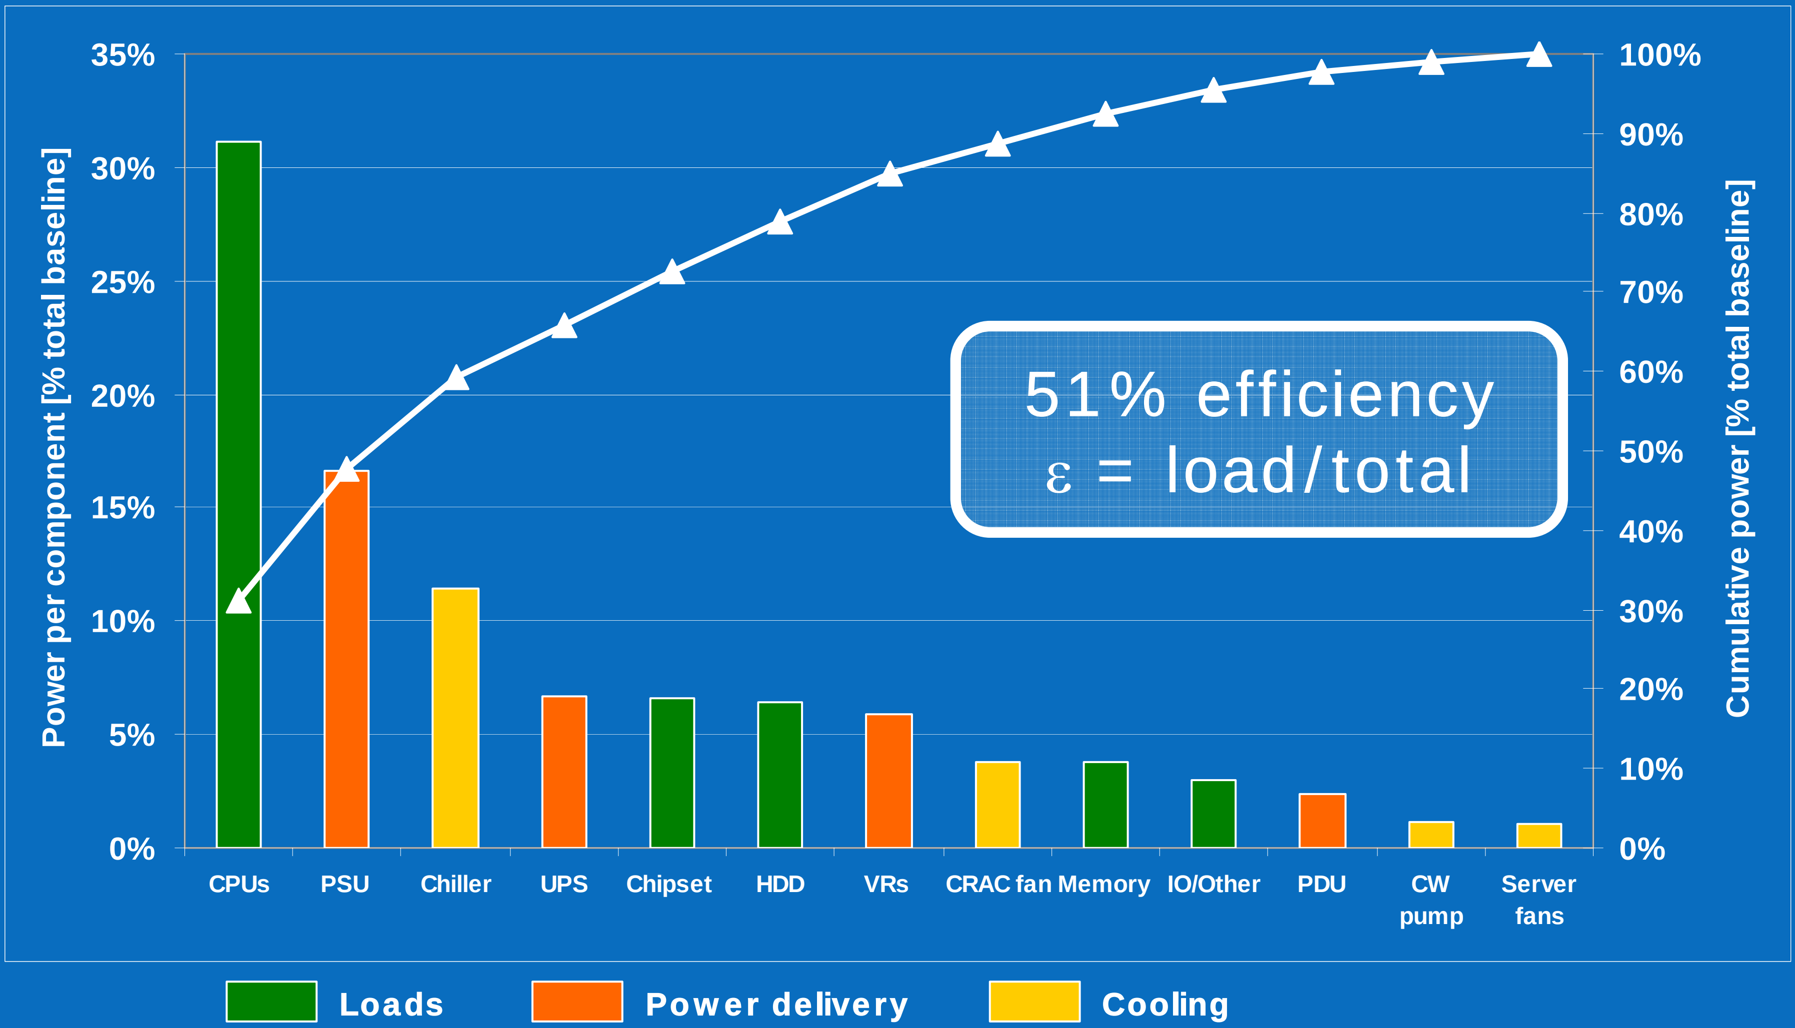

Opportunities

Pareto shows results of each area/component in the Data Center

Compute uses *(CPU, memory, IO, HDD, chipset)*

CPUs use the most power but still only 20% of the total power budget

Memory power increasing as processor capabilities improve

Power Delivery and Conversion uses *(PSU, UPS, VR, PDU)*

AC/DC, DC/AC, and DC/DC power conversions use large amounts of energy

Cooling uses *(Chiller, System Fans, CRAC Fans, CW Pumps)*

Chiller is the top contributor in cooling

Data Center airflow well down on the list

Strained Silicon Transistors

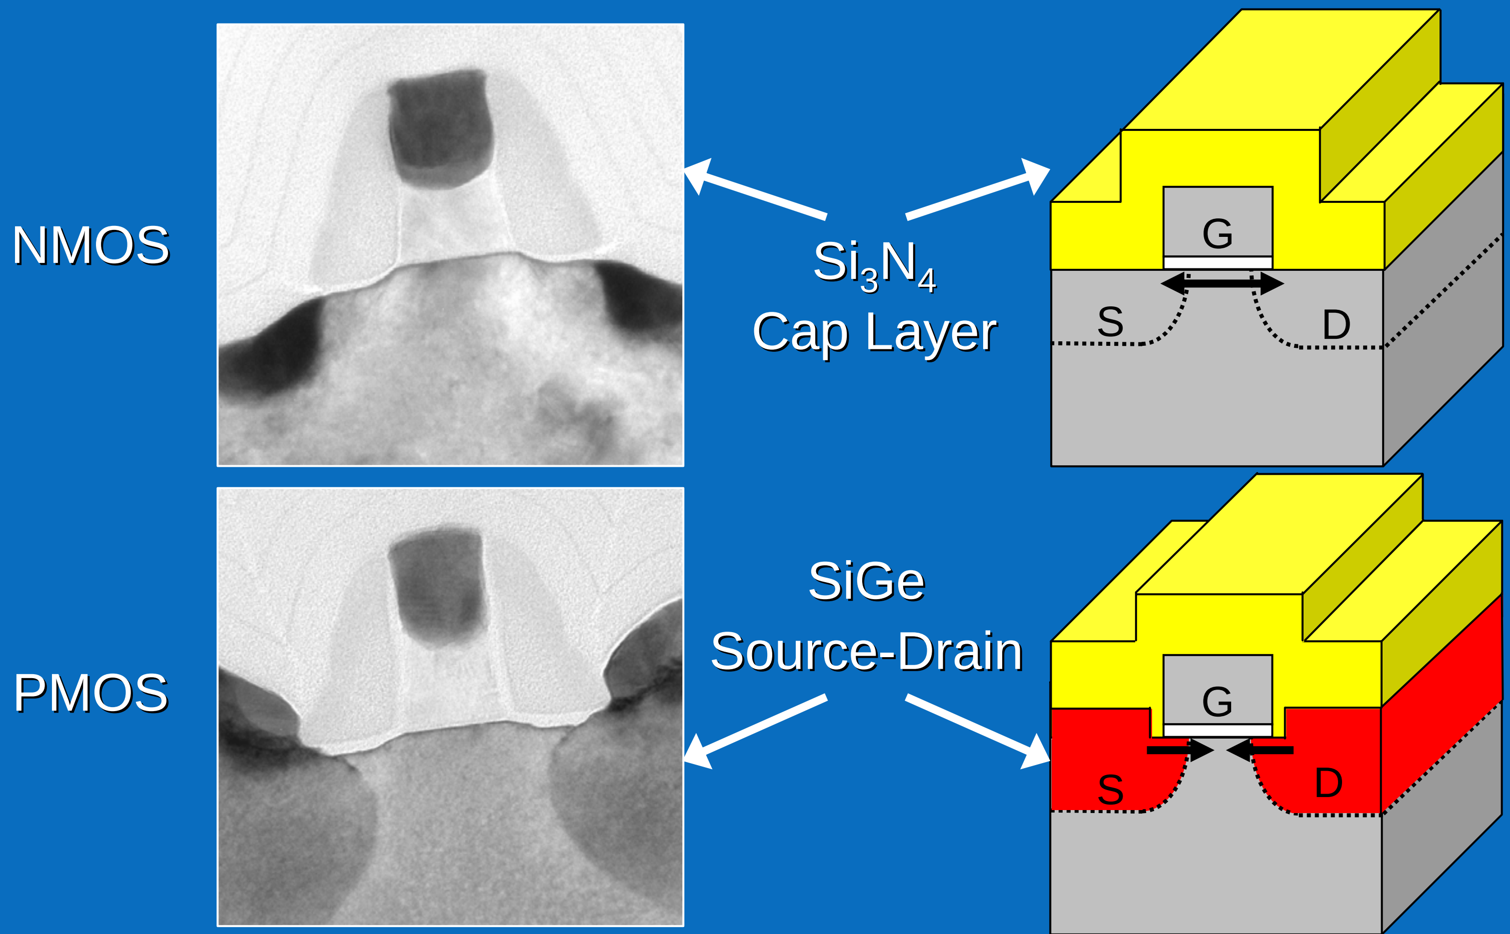

2ND generation strained silicon for enhanced performance

35 nm gate length   1.2 nm gate oxide   NiSi for low resistance

Transistor Performance vs. Leakage

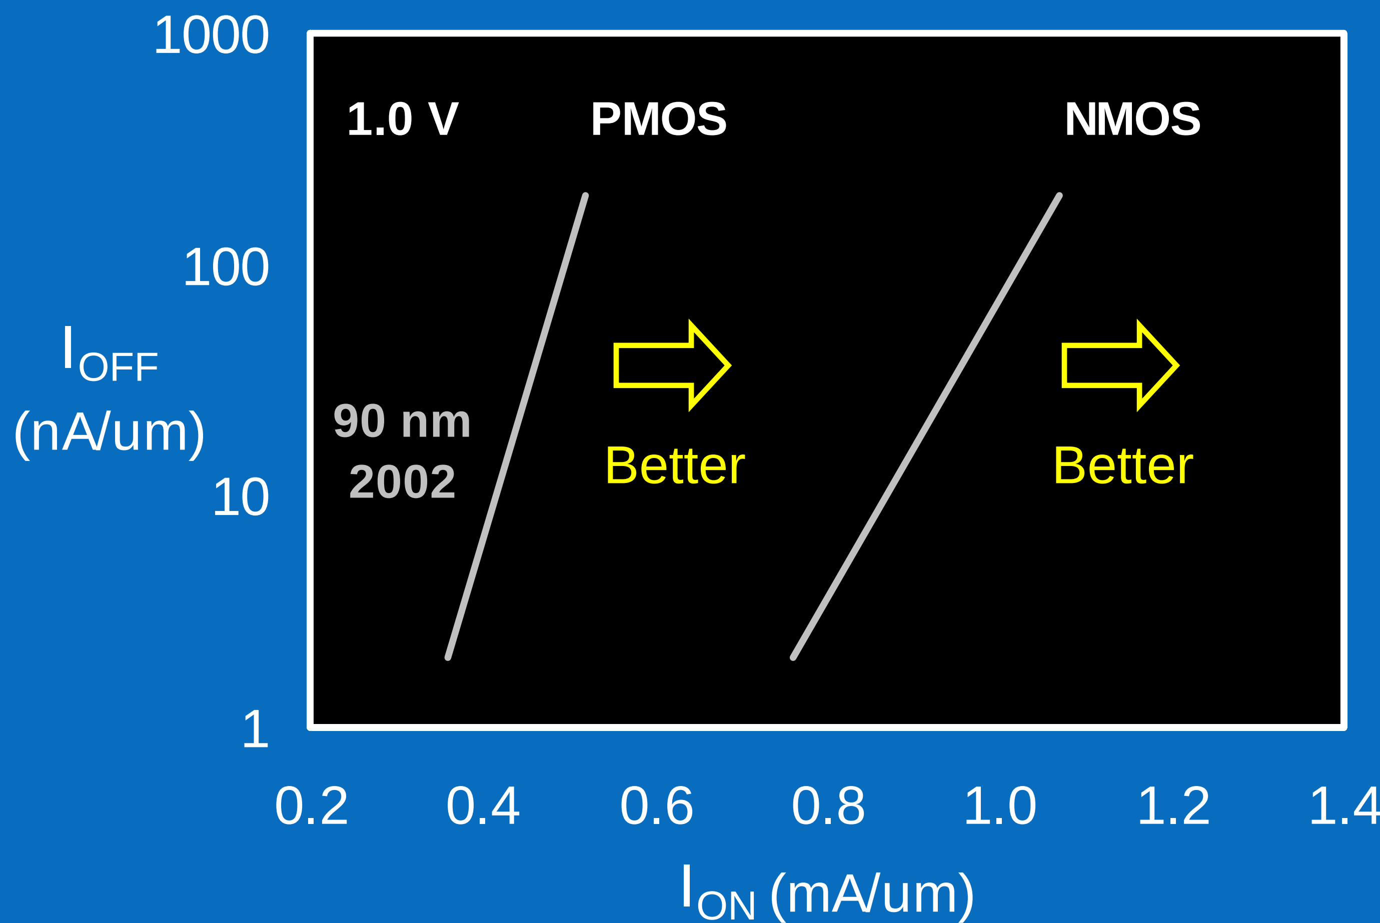

Improved transistors provide increased drive current (I_{ON}) at constant leakage current (I_{OFF})

Improved Transistor Performance

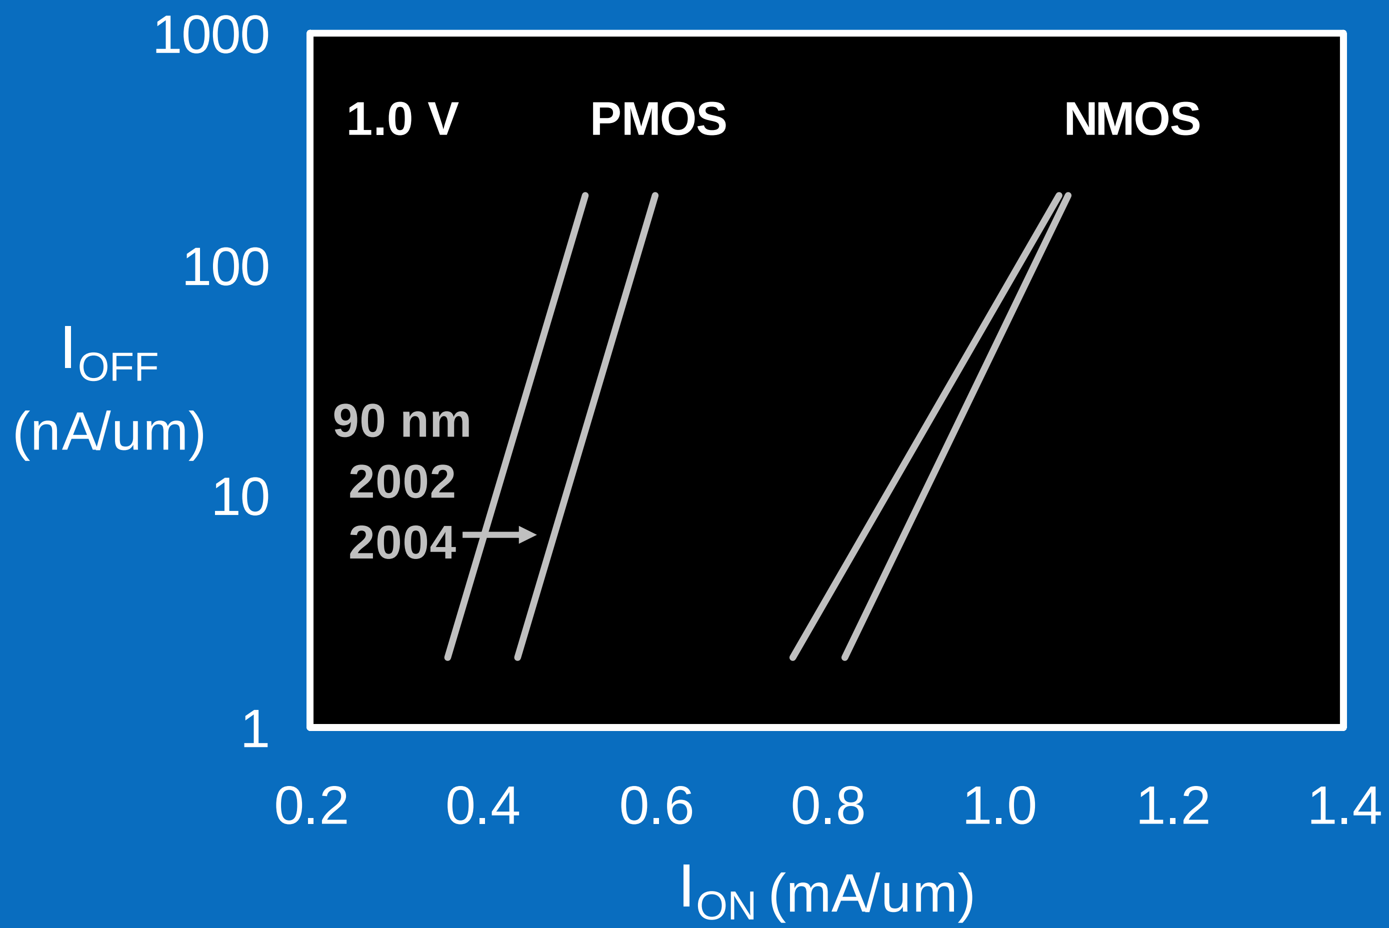

90 nm transistors have continued to improve

Improved Transistor Performance

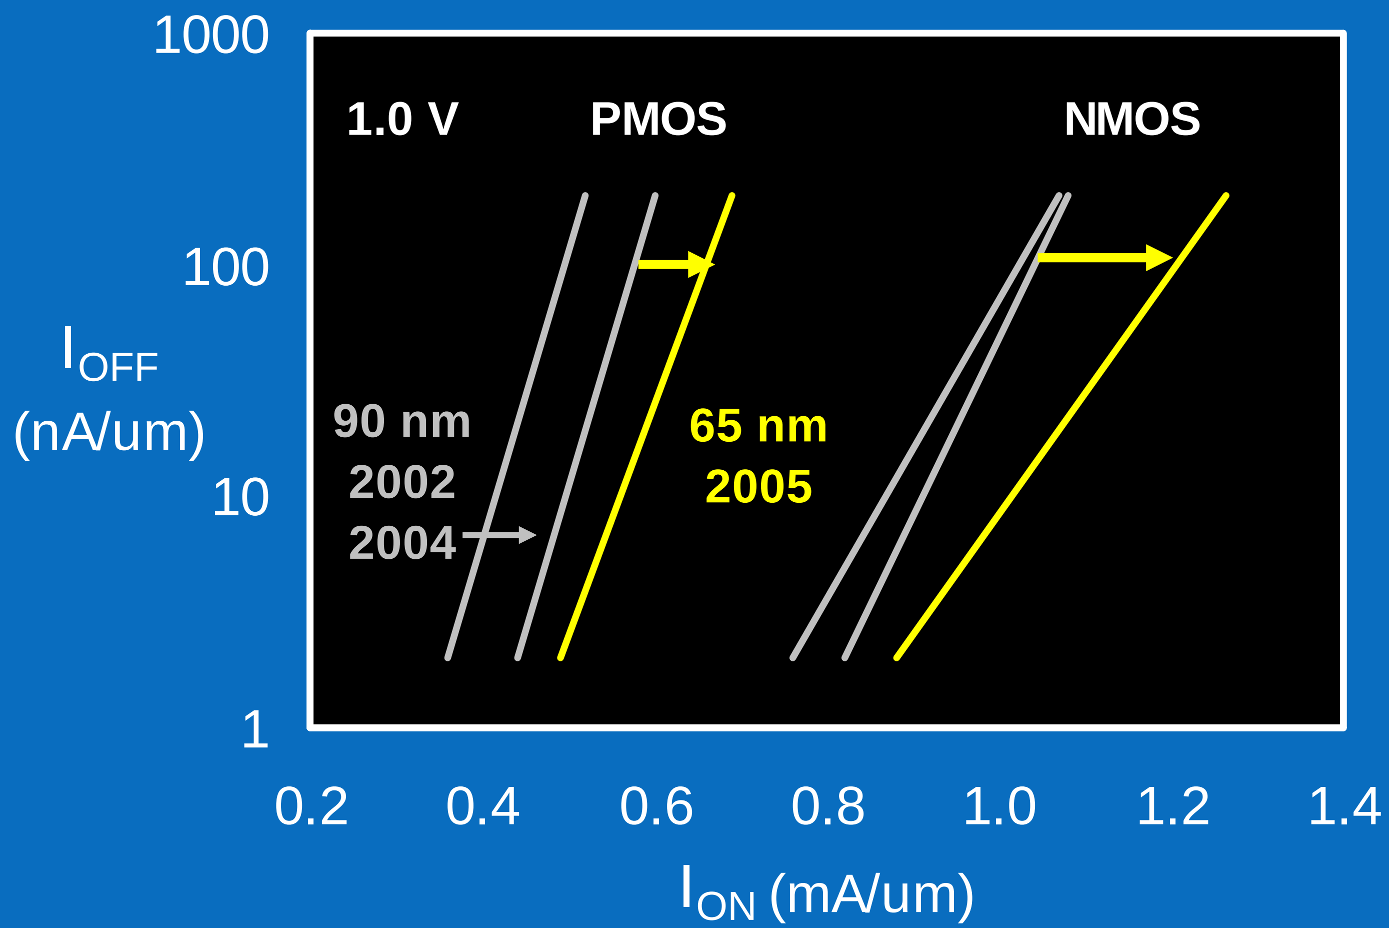

65 nm transistors increase drive current ~15% with enhanced strain

Reduced Transistor Leakage

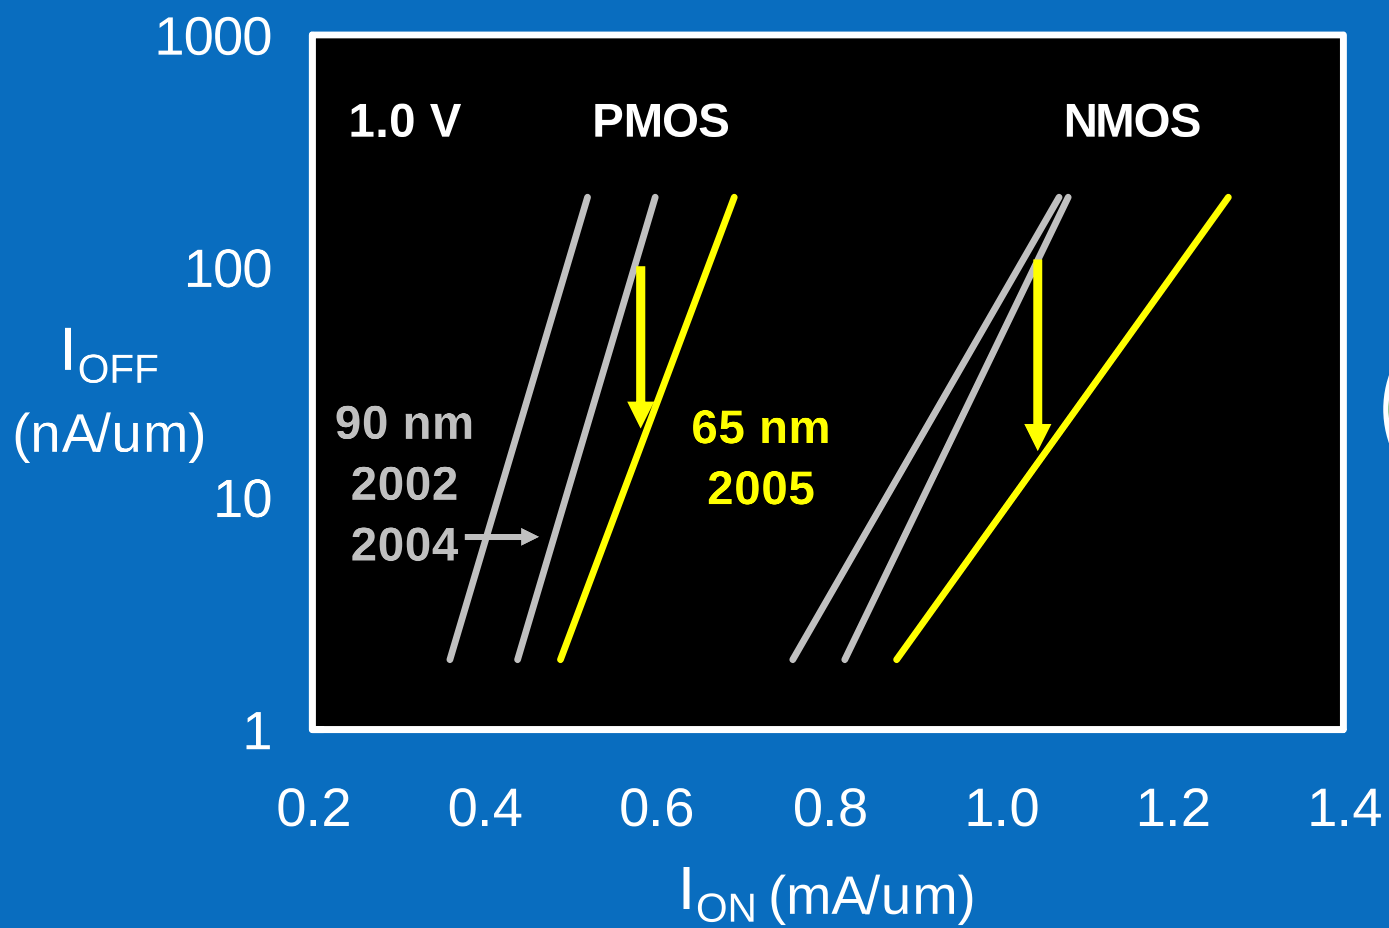

Power Saving Feature

65 nm transistors can alternatively provide ~5x leakage reduction

Silicon Opportunities

Strained silicon is just one of many technologies on-going to reduce chip power. Material changes, architectural improvements, transistor designs, and more all are being used and focused on lower power.

Intel® Xeon™ performance optimized processor details:

	Processor	Technology	Cores	TDP
2005	Irwindale	90nm	1	110 W
2006 H2	Woodcrest	65nm	2	80W

Power Delivery Opportunities - Server

Today

Intel motherboard VR efficiency improved 5% over the last 7 years, with 5x increase in current.

PSU is main culprit, as shown in pareto

- SSI Industry specifications have been introduced to encourage more efficient power supplies
- 80 PLUS program covers 1U server power supplies but implementation is limited
- Intel supports EPA's goals with proposed revision of Energy Star specification and is active in providing feedback
- Technology exists for 90% efficient power supplies – cost remains barrier to use, *despite* ROI < 1 year

Pathfinding

Intel continues to work at improving power delivery efficiency

Power Delivery Opportunities - Server

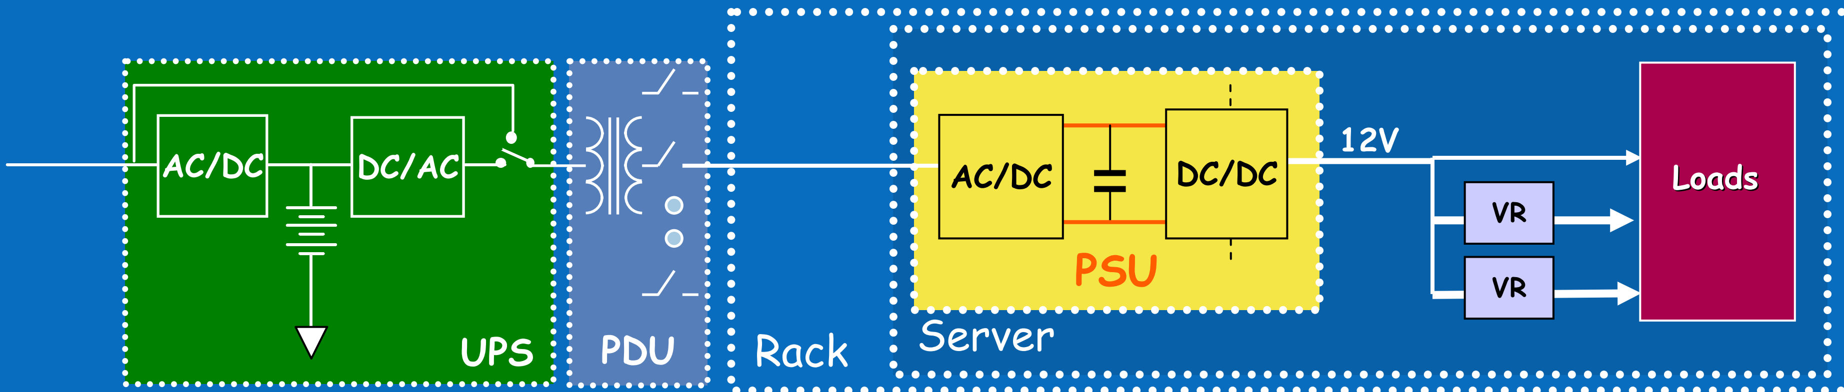

Eliminate redundant regulation stages

efficiency 86% → 96%

Use more efficient DCX

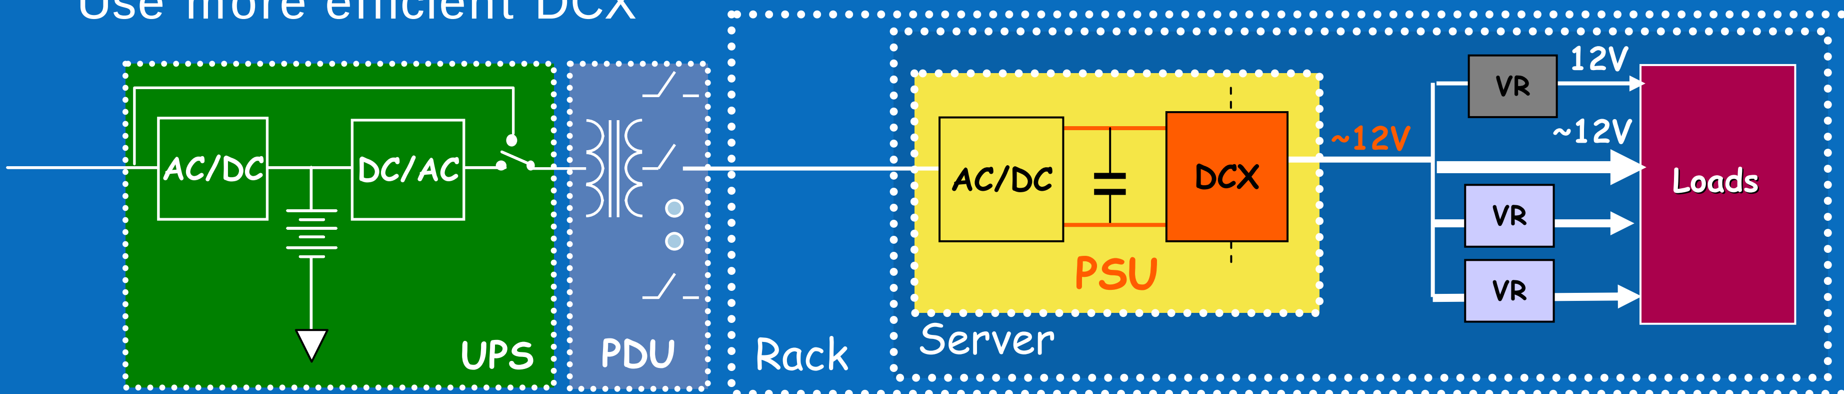

DCX ⇒ Unregulated DC/DC Converter; DC Transformer

Power Delivery Opportunities - Room

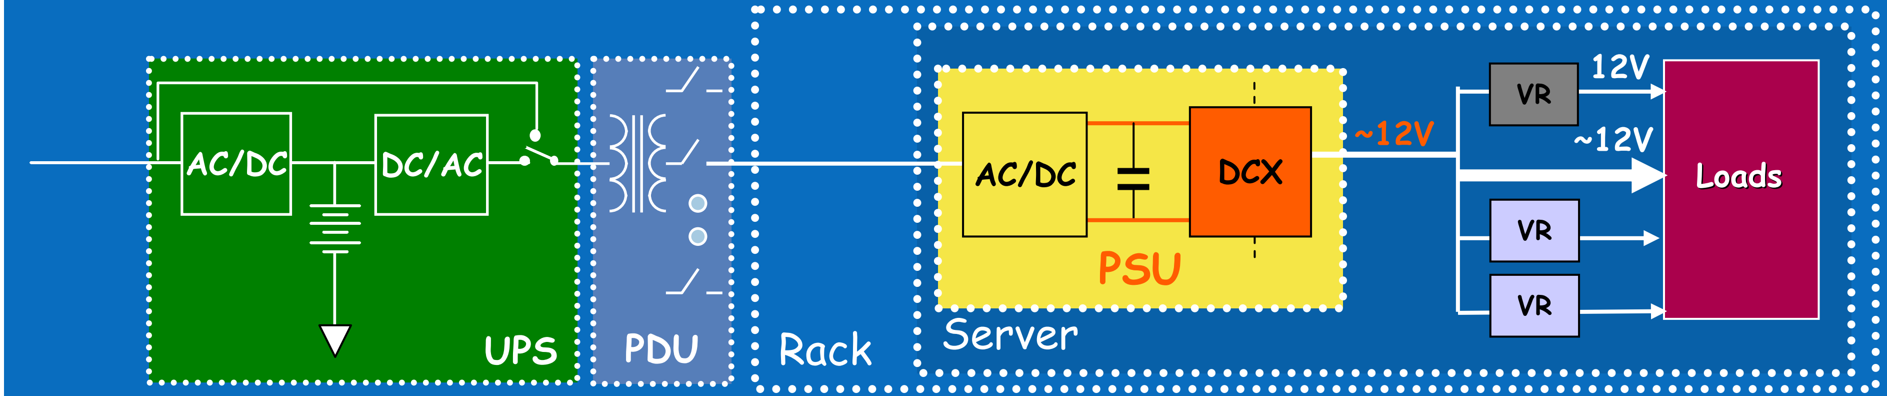

Move AC/DC to rack to reduce heat load and PSU volume in servers

Collaboration with Lawrence Berkeley Nat'l Lab & other industry partners

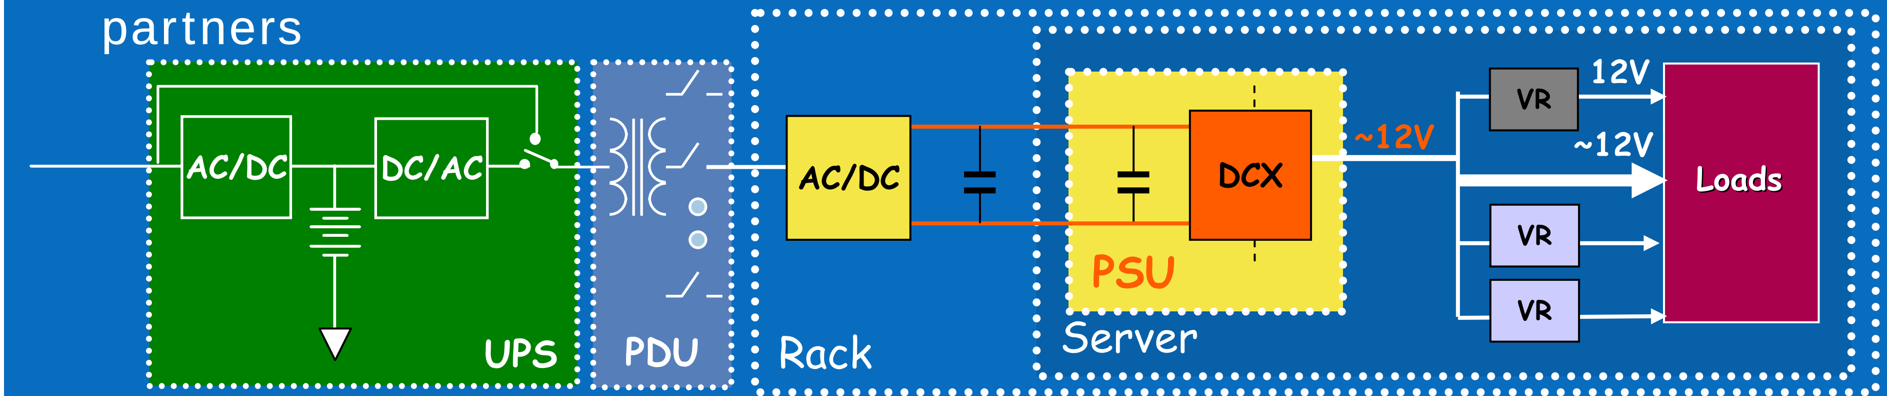

efficiency 93% → 97%

Power Delivery Opportunities - Room

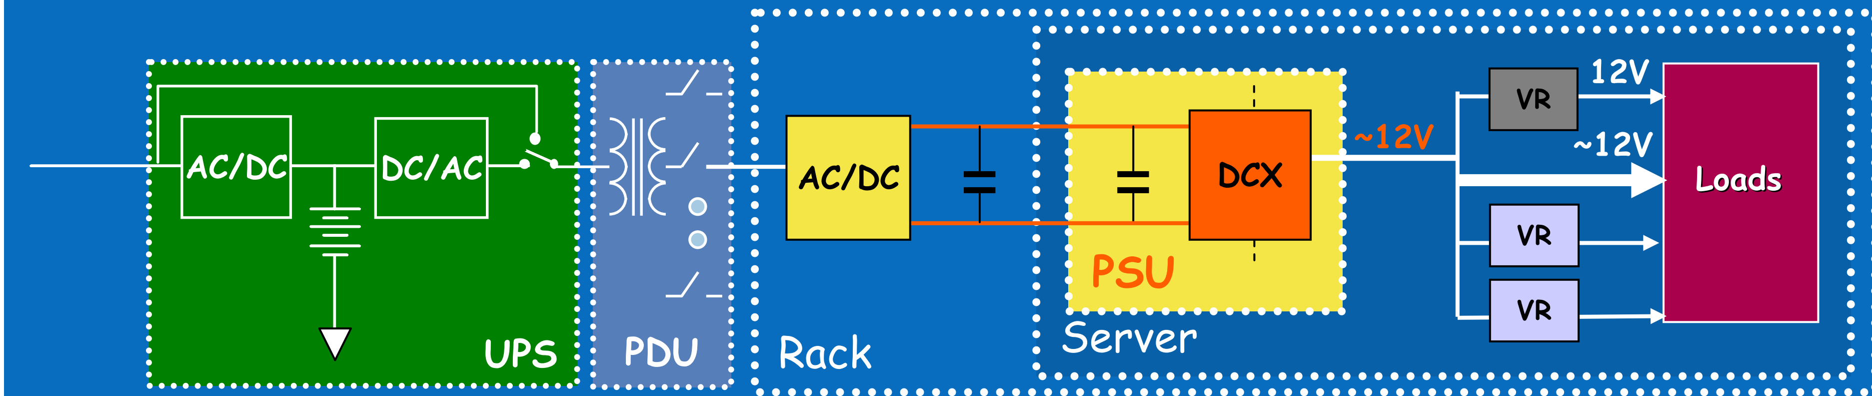

High Voltage DC (~380V) in data center increases efficiency through reduction in conversion stages

Collaboration with Lawrence Berkeley Nat'l Lab & other industry partners

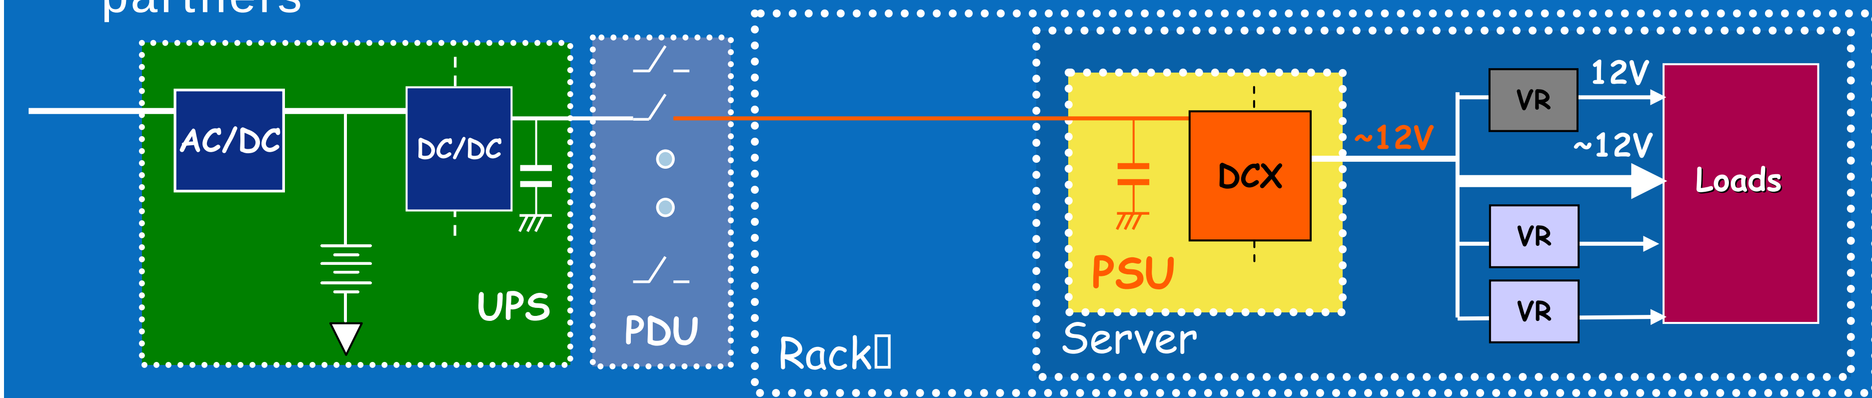

efficiency 93% → 96%

IR photograph of data center rack inlets

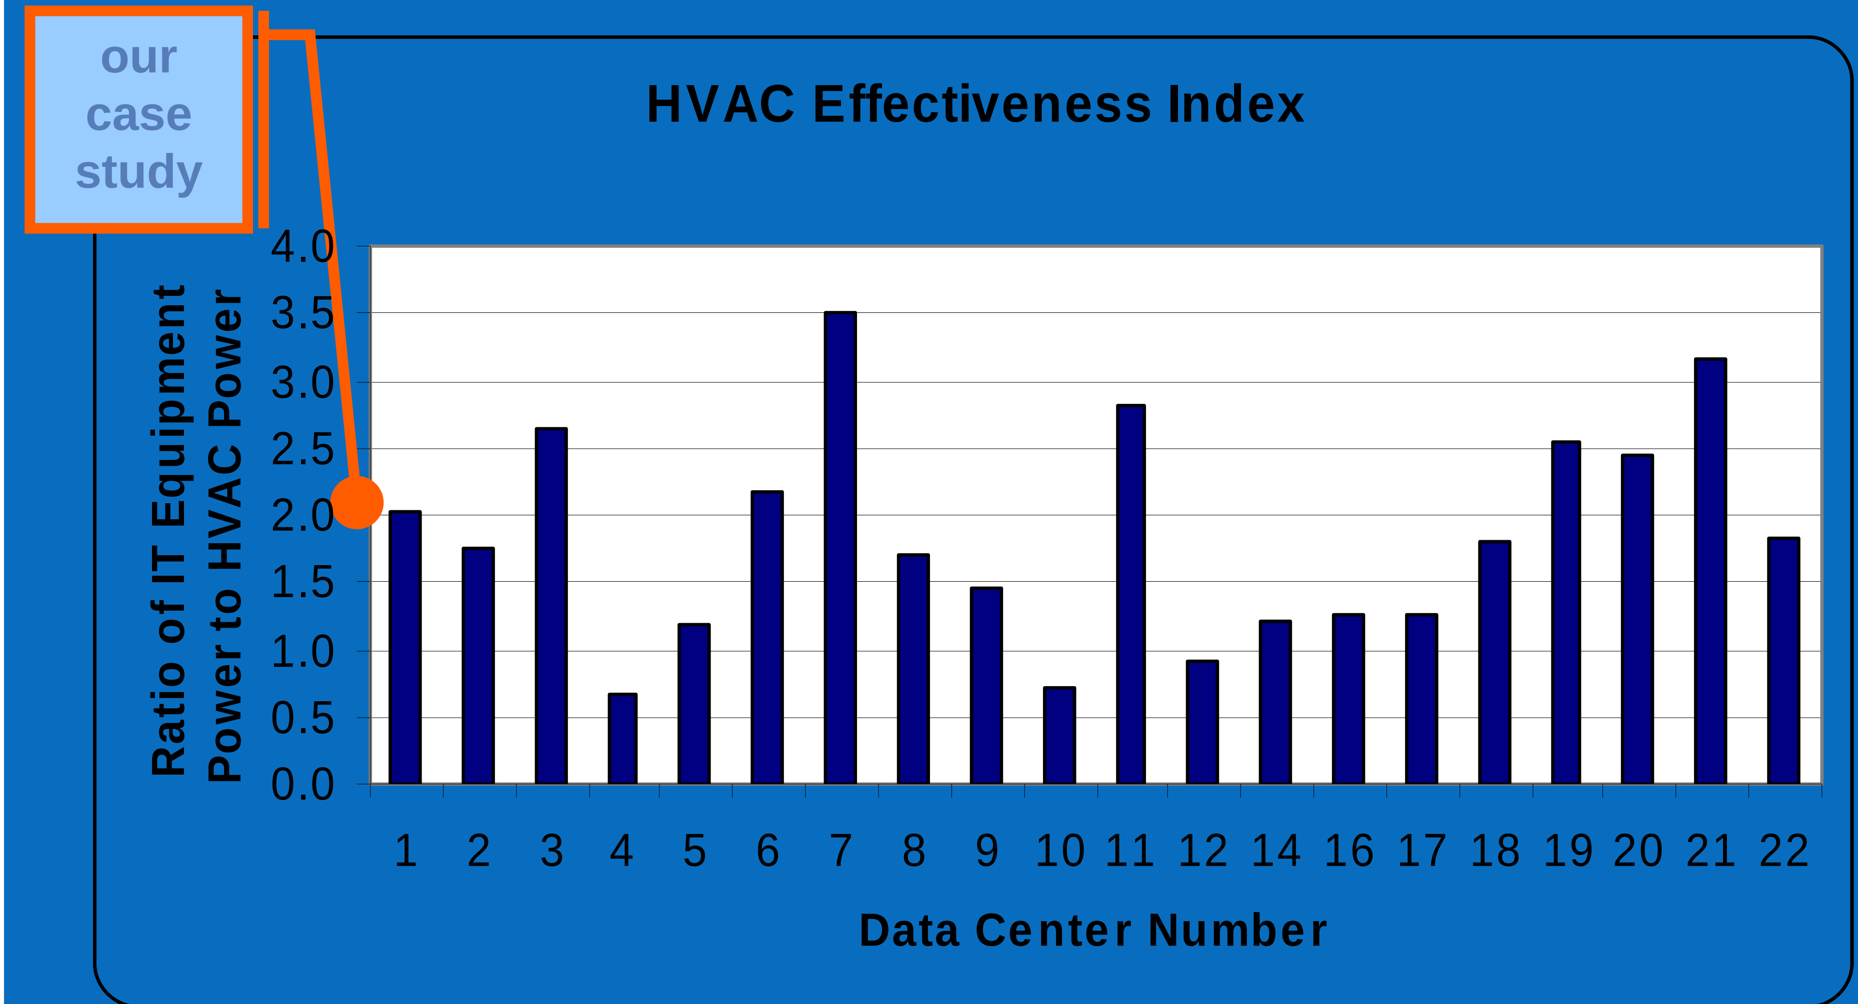

Slide adapted from LBNL DC energy study, W. Tschudi

Cooling Review

Chiller was #3 on the pareto (11% of the total budget)

Installation used air-cooled chillers. Not optimum efficiency. Liquid cooled chillers (e.g. cooling towers) more efficient. Additional techniques such as air-side economizers and wet-side economizers can improve efficiency even further.

Chiller efficiency is a campus-wide issue, difficult to drive from the DC

Capital vs Operational trade-off during initial construction, rarely a good candidate for retrofit efficiency improvements.

CRAC fans were #8 on the pareto (~4% of the total budget)

Airflow management is key to a healthy Data Center. If done poorly the affect is on server performance, more than on TCO.

High Density DC vs Low Density DC

Same Chiller and CW pumps for both, main change is airflow at a higher static pressure. Airflow only 4% of total at High Density!! Operational cooling costs do NOT drive us to low density DCs.

Possible results – implementing some of the discussed technologies

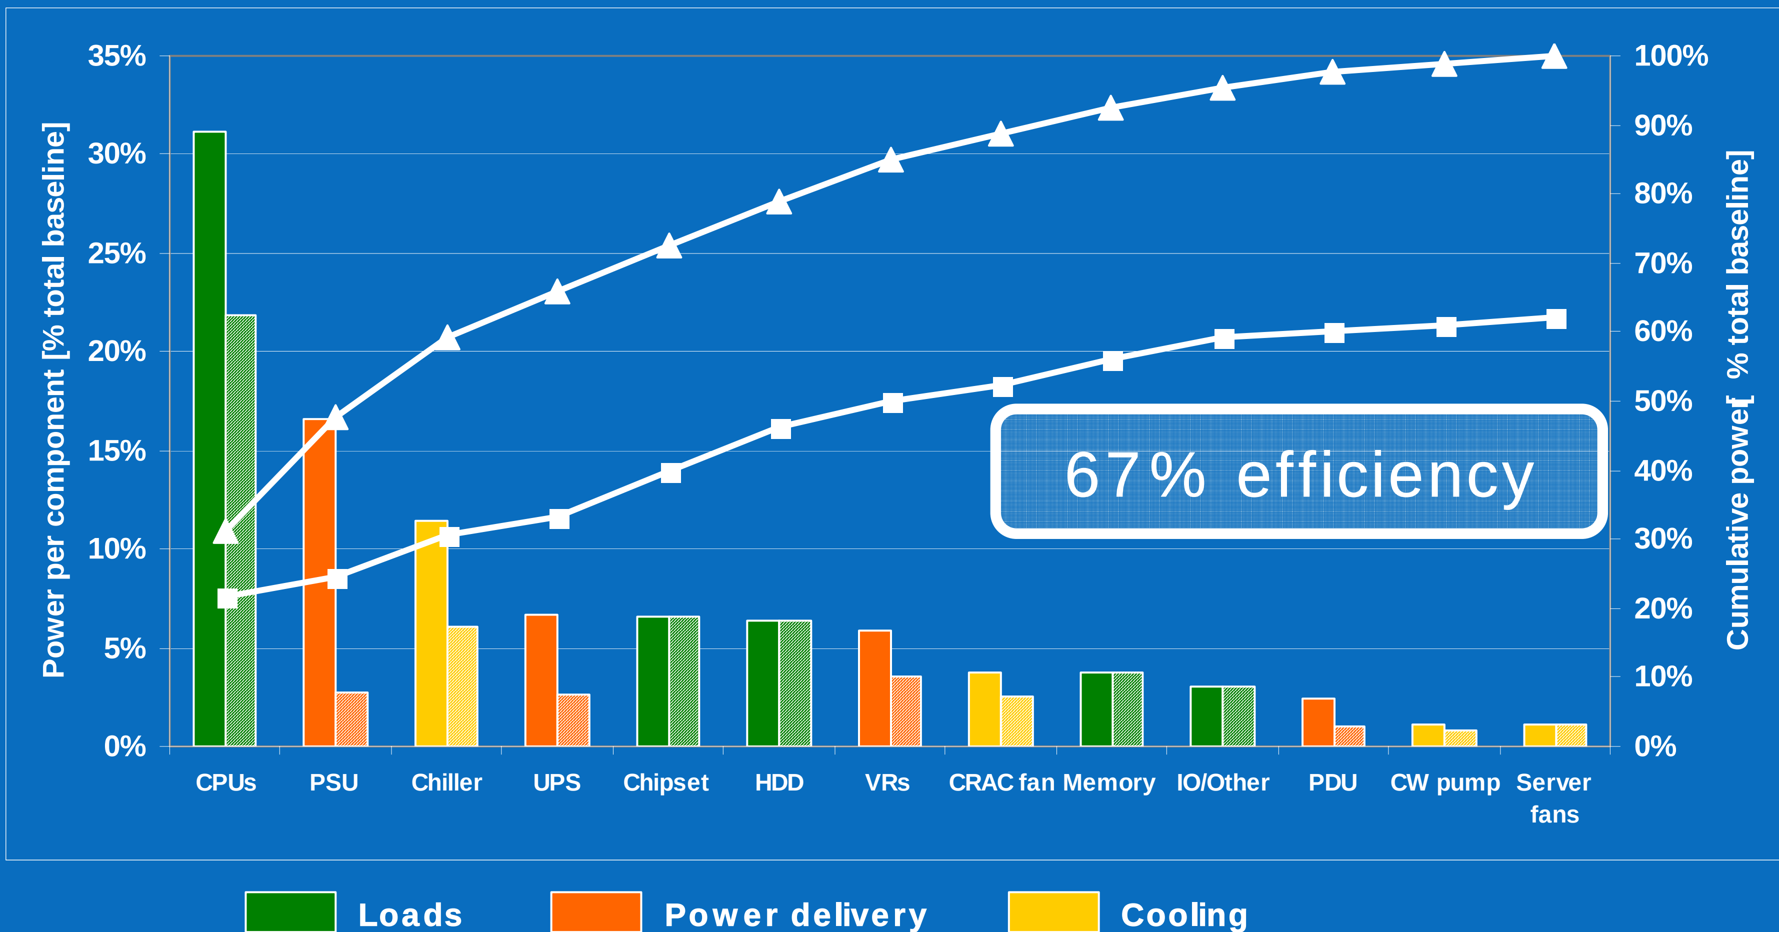

Summary and what's needed

Opportunities

Silicon technology continues to improve, providing more performance/watt

More efficient cooling is available, needs end-user adoption

Power delivery architecture improvements in development but further out and require supplier support and end-user demand to accomplish

Efficiency and Metrics

Showed the possibility to move this DC from 51% to 67%

Just changing to a dual-core lower power processor, ε goes up to 52%

Metric lacks compute-performance aspect and is essentially broken!

We need metrics of performance/watt$_{server}$ and performance/watt$_{total\ DC}$

Know your Data Center's characteristics! Learn where the real opportunities are. Attack the big power users first.